CRD report

News from the LBNL Computational Research Division

November 2004

LBNL Establishes Berkeley Institute for Performance Studies

Alan Laub, former director of DOE's SciDAC program, once said that in the field of high performance computing, "peak performance" was defined as the speed at which "the manufacturer guarantees that you can't compute faster than that." Although peak performance figures make for good marketing, they don't provide much insight into actual performance.

To rectify this, for the past eight years Berkeley Lab has been developing new tools and techniques for more accurately assessing the performance of high performance computers, especially when it comes to running real-world scientific applications.

In November, many of these projects in performance characterization, modeling and benchmarking for supercomputers were brought together to comprise the Berkeley Institute for Performance Studies. Known as BIPS, this umbrella organization will be led by Kathy Yelick (see sidebar) and encompasses the following research activities at LBNL and UC Berkeley:

The Performance Evaluation Research Center (PERC), directed by David Bailey, is one of seven SciDAC Integrated Software Infrastructure Centers (ISICs). PERC involves approximately 25 researchers at eight centers (four labs and four universities). The goal of PERC is to develop a science for understanding performance of

(continued on page 2)

CRD Report

CRD Report is published every other month, highlighting recent achievements by staff members in Berkeley Lab's Computational Research Division. Distributed via email and posted on the Web at http://crd.lbl.gov/DOEresources, CRD Report may be freely distributed. CRD Report is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

Crystallization in Silico

Figuring out the structures of biological macromolecules

When Francis Crick and James Watson deciphered the structure of DNA in 1953, X-ray crystallography became famous; key to their success was crystallography of DNA done by Rosalind Franklin in the laboratory of Maurice Wilkins. X-ray crystallography has long since become the workhorse for structural studies of big biological molecules, including most of the many thousands of proteins whose structures have been solved in the last half century.

Crystallizing biological molecules is tricky, however. Some proteins and other macromolecules can't be crystallized at all; those that can must first be painstakingly purified. And a molecule's shape as part of a crystal, a highly artificial state, may significantly differ from its shape (or shapes) in the warm, aqueous environment of a living cell.

Enter single-particle electron cryomicroscopy (cryo-EM). Bob Glaeser, a member of Berkeley Lab's Life Sciences and Physical Biosciences Divisions and a professor of bio-

(continued on page 3)

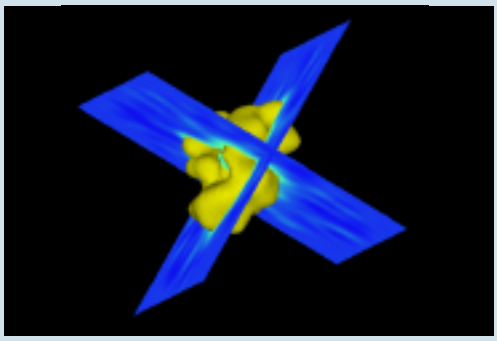

In single-particle reconstruction, 3-D structures of macromolecules or molecular assembly are obtained by merging 2-D projection images of isolated (single) particles with random and unknown orientations. This figure shows an isosurface rendering of the 3-D structure of a multisubunit transcription factor (TFIID) and the relative orientations of two of the projection images. The Fourier transforms of these projections intersect along a common line in the Fourier space of the 3-D structure. The orientations of the projection images are unknown in advance and must be computed during the reconstruction process.

DataMover Reaches Milestone with Automated Transfer of 18,000 Files in a Single Request from Brookhaven to Berkeley

Amidst the hype and hoopla at the recent SC2004 conference in Pittsburgh, CRD's Scientific Data Management Research Group demonstrated the robustness of the group's DataMover by putting the application through its workaday paces. In doing so, the group reached a milestone when, in a single request, 17,870 data files were moved seamlessly from Brookhaven National Lab to LBNL.

What made the transfer significant was that it was steered by LBNL's Eric Hjort from the conference in Pittsburgh, and that the number of files moved was the highest ever. But it was just another day of moving data for Hjort, who oversees the moving of files generated at Brookhaven's STAR experiment to the High Performance Storage System (HPSS) at the NERSC Center. DataMover automates all aspects of the transfer once the user determines which directory and all its files are to be moved and where they will be moved to.

Once the application starts, DataMover communicates with the source and target hierarchical resource managers (HRMs) at BNL and NERSC. The HRMs are Grid middleware components developed by the Scientific Data Management Research Group that manage staging and archiving of files from/to HPSS. The DataMover extracts the directory structure from the source HPSS through HRM, generates the corresponding directory structure at the target HPSS through HRM, and puts the list of requested files in the target hierarchical resource manager (HRM). The target HRM then contacts the HRM at the

(continued on page 4)

BIPS Established *(continued from page 1)*

scientific applications on high-end computer systems, and develop engineering strategies for improving performance on these systems. The project is integrating several active efforts in the high performance computing community and is forging alliances with application scientists working on DOE Office of Science missions to ensure that the resulting techniques and tools are truly useful. For detailed information about PERC, go to http://perc.nersc.gov/main.htm.

The Berkeley Benchmarking and Optimization Group (BeBOP) is led by Kathy Yelick and James Demmel of UC Berkeley, with substantial participation by Berkeley graduate and undergraduate students. Their research areas include:

- the interaction between application software, compilers, and hardware
- managing trade-offs among the various measures of performance, such as speed, accuracy, power, storage
- automating the performance tuning process, starting with the computational kernels which dominate application performance in scientific computing and information retrieval
- performance modeling and evaluation of future computer architectures.

The BeBOP Web site can be found at http://bebop.cs.berkeley.edu/.

BeBOP works closely with the UCB LAPACK/ScaLAPACK project, which focuses on new algorithms for numerical linear algebra and new, more efficient implementations of linear algebra software.

Berkeley Lab's architecture evaluation research project, led by Leonid Oliker and Yelick, is conducted by staff from CRD and the NERSC Center Division, as well as collaborators from other institutions. They evaluate emerging architectures, such as processor-in- memory and stream processing, and develop adaptable "probes" to isolate performance-limiting features of architectures. They conducted the first in-depth analysis of state-of-the-art parallel vector architectures, running benchmark studies on the Japanese Earth Simulator System (ESS) and comparison runs on Cray's X1 system. Results on the ESS demonstrated 23 times faster performance than the IBM Power3 in a node-to-node comparison. (See the September issue of CRD Report for more information on this work.)

NERSC's benchmarking and performance optimization project is carried out by NERSC staff with expertise in performance analysis. They developed the Effective System Performance (ESP) benchmark to measure system-level efficiency and the Sustained System Performance (SSP) benchmark to measure overall system application throughput. SSP resulted in a 30 percent increase in the Seaborg system's capability and is now used in several non-DOE procurements. This team also accelerated several SciDAC application programs running on Seaborg. Read more about ESP at http://www.nersc.gov/projects/esp.php.

KATHY YELICK NAMED LEADER OF BIPS AND CRD S FUTURE TECHNOLOGIES GROUP

Kathy Yelick

Kathy Yelick, a professor of computer science at UC Berkeley with a joint appointment in the Computational Research Division, has been named to lead for the newly established Berkeley Institute of Performance Studies (BIPS). She will also be leading CRD's Future Technologies Group (FTG). Yelick's appointment, which includes a leave of absence from her teaching position, officially takes effect Jan. 1, 2005.

The main goal of Yelick's research is to develop techniques for obtaining high performance on a wide range of computational platforms and to ease the programming effort required to obtain improved performance. She is perhaps best known for her efforts in global address space (GAS) languages, which attempt to present the programmer with a shared memory model for parallel programming. These efforts have led to the design of Unified Parallel C (UPC), which merged some of the ideas from three shared address space dialects of C: Split-C, AC (from IDA), and PCP (from LLNL). In recent years, UPC has gained recognition as an alternative to message passing programming for large-scale machines. Compaq, Sun, Cray, HP, and SGI are implementing UPC, and Yelick is currently leading a large effort at LBNL to implement UPC on Linux clusters and IBM machines and to develop new optimizations. The language provides a uniform programming model for both shared and distributed memory hardware. Read more at http://upc.lbl.gov/. She has also worked on other global address space languages such as Titanium, which is based on Java.

Yelick has also done some notable work on single processor optimizations including techniques for automatically optimizing sparse matrix algorithms for memory hierarchies. These efforts are part of an NSF-funded project called BeBOP (Berkeley Benchmarking and Optimization) that is working on methods to take advantage of special structure such as symmetry and triangular solves.

Another area that Yelick has worked on that has led to very interesting results is her research on architectures for memory-intensive applications and in particular the use of mixed logic and DRAM, which avoids the off-chip accesses to DRAM, thereby gaining bandwidth, while lowering latency and energy consumption. In the IRAM project, a joint effort with David Patterson, she developed an architecture to take advantage of this technology. The IRAM processor is a single chip system designed for low power and high performance on multimedia applications and achieves an estimated 6.4 GOP/s in a two-watt design. The IRAM architecture is based on vector instructions, historically reserved for expensive vector supercomputers designed for large-scale scientific and engineering applications.

Yelick earned her bachelor's (1985), master's (1985), and Ph.D. (1991) degrees in electrical engineering and computer science from the Massachusetts Institute of Technology. Her research interests include parallel computing, memory hierarchy optimizations, programming languages and compilers. You can read her UC Berkeley Web page at http://www.cs.berkeley.edu/~yelick/.

Crystallization in Silico

(continued from page 1)

chemistry and molecular biology at UC Berkeley, explains that instead of trying to build a crystal in which vast numbers of biological macromolecules assume regular spacing and orientation, there's a different approach.

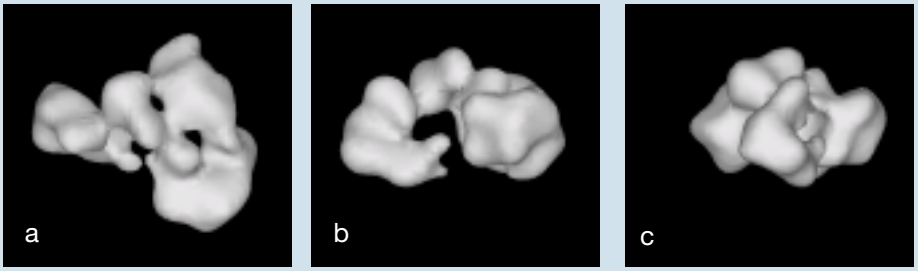

The isosurfaces of the TFIID structure from three different viewing angles. These surface renderings are generated by Vis5D. The leftmost view (a) is the top view; (b) is the front view obtained by rotating (a) by 90 degrees around the horizontal axis; and (c) is obtained by rotating (b) by another 90 degrees around the vertical axis.

With cryo-EM, Glaeser says, "you can put a microliter of a reasonably pure sample in aqueous solution onto a carbon support film, then plunge it into ethane at liquid-nitrogen temperature" — which freezes the solution so rapidly that the water in it becomes vitreous, or glassy. The frozen sample is then put under the electron microscope to create two-dimensional images of thousands of randomly oriented "single particles" of the macromolecule.

"To see the structure in 3-D, you have to merge the data from all these individual images, whose orientations are not known," says Glaeser. "However, once these images have been aligned computationally, in a known orientation relative to one another, you have effectively constructed an artificial crystal in the computer" — what Glaeser calls "crystallization in silico."

There are a couple of catches, he says. "One is that the electron beam can do a lot of damage, and imaging with a safe but weak exposure results in noisy images." The low ratio of signal to noise complicates the process of identifying 2-D images from the micrographs suitable for constructing the 3-D model, and of eliminating spurious data during the calculation.

Another, more fundamental catch is the number of calculations required. "About 100,000 particles are enough to pick out an alpha helix," Glaeser says, "but if you want atomic resolution, good enough to resolve a polypeptide chain, you'll need a million particles or more." Glaeser says that to achieve 3-angstrom resolution by analyzing a million particles using the most straightforward methods available today would require on the order of 10^{24} arithmetic operations. "Today's best machines would take 10^{10} seconds to run the calculation," he says — almost 20,000 years.

Determined to overcome these limitations of the cryo-EM technique, Glaeser approached members of the Lab's Computational Research Division (CRD) for help in improving mathematical approaches to constructing 3-D images from single particles. Initially this work was supported by Laboratory Directed Research and Development funding, but the project soon attracted the attention of the National Institutes of Health. In 2003 NIH launched an ongoing multidisciplinary, multi-institutional Program Project titled "Technology Development for High Resolution Electron Microscopy." Besides Glaeser, the project's seven directors include two of the early developers of single-particle software, Joachim Frank of the New York State Department of Health's Wadsworth Center in Albany, and Pawel Penczek of the University of Texas's Houston Medical Center.

One major effort lies in developing new computational approaches and improving algorithms to determine the orientation of the 2-D images and continually refine the construction of the 3-D model from these. Esmond Ng, head of the Scientific Computing Group in CRD, enlisted group member Chao Yang to help meet the challenge. They early on set out to learn the basics of single-particle cryo-EM through collaboration with co-PI Penczek, Ken Downing of the Lab's Life Sciences Division, and Eva Nogales of the Life Sciences and Physical Biosciences Divisions, an associate professor of molecular and cell biology at UC Berkeley.

Says Ng, "Once we understood the problem and the issues the microscopists were facing in reconstructing 3-D models from a selection of randomly oriented 2-D projections, we sought new ways of formulating the problem mathematically. We realized there were computational tools for tackling some formulations already in existence. The tools aren't new, but structural biologists don't know about them or haven't used them. So both parties had something to offer the other; that's the beauty of this collaboration."

Yang characterizes one approach as "top down — describing the general problem and looking for the best numerical solution, the best algorithm. The experimentalists come from the bottom up, coping with specifics. Now we are converging, working toward a robust algorithm that can handle peculiar problems, like noisy data in cryo-EM."

In an article to be published in the *Journal of Structural Biology*, Yang, Ng, and Penczek describe an algorithm for simultaneously refining the 3-D model while tightening the parameters for the orientation of the individual 2-D projections used to reconstruct the model. The method is faster, more efficient, and more accurate than any of its predecessors.

Because they are projections, the electron-microscope's many images of identical proteins — quick-frozen from solution, on carbon film — look different from one another, just as shadows of identical pasta tubes on a flat piece of paper would show two concentric circles seen along the axis, concentric ellipses seen at an angle, and a dark band seen from the side. Without knowing the exact orientations of these views, however, one might not be able to tell if the pasta tubes were cut straight across like macaroni or slanted like cannelloni.

In the case of proteins — with shapes generally more complex than pasta! — the first task is to select enough good-quality images. Once enough projections have been chosen, they can be grouped according to their apparent orientations on the carbon film and averaged, in order to improve the signal-to-noise ratio. From these groups, preliminary 3-D models are constructed.

The first model is only an educated guess of the real final shape. This model becomes increasingly more accurate, however, as the

(continued on page 4)

Single-Particle Reconstruction

(continued from page 3)

orientations of the selected particles are continually corrected and the model refined.

Usually the refinement of orientations and the subsequent refinement of the model itself are done as separate steps in "real space." In a leading method known as projection matching, developed by Penczek, the individual particles are reoriented in a way that corresponds best to the current model; then the model is refined to better fit the sum of the projections, and the process is repeated until no further improvement is possible. With this method, unfortunately, mistakes introduced into the model at any stage are likely to persist.

Alternatively, the data can be mathematically transformed so that particle orientations are corrected all together, with the advantage that the final 3-D model need be calculated only once. But in these approaches the model calculation isn't based on corrections made in real space, and the mathematical transformations themselves introduce uncertainties and possible errors.

Yang, Ng, and Penczek's new method simultaneously optimizes particle orientation and model refinement in real space. Unlike projection matching, it does not correct orientations by comparison with the model.

Says Yang, "If you know you don't have an optimum 3-D structure, do you really want to try all that hard to match it? Instead, our approach uses derivative information to search for the minimum difference between particle orientations and various model configurations in a cubic grid. All you need is a search direction; you compute on the fly."

Technically, the problem is formulated as an optimization problem and solved using the limited-memory Broyden-Fletcher-Goldfarb-Shannon (BFGS) algorithm; the mathematically inclined will find details in the authors' Journal of Structural Biology paper, referred to below. It's a computational method that has been applied in many scientific fields; its application to single-particle cryo-EM, using the supercomputers at NERSC, is a significant step forward, offering a rapid, robust way — one in which the very nature of the calculation tends to eliminate noise and bad data — of achieving dependable structures at medium resolution.

In the quest to achieve atomic-resolution structures of large biological molecules in solution, many challenges remain. They include finding the best software designs for different computer architectures; finding ways to handle the data from million-particle collections with fewer operations and faster calculations; and — from the standpoint of the biologist, one of the most desirable goals of all — the ability to study the same protein in different conformational states. The last is a goal that crystallography renders unattainable by its very process, but one that "crystallization in silico" brings closer to realization.

Additional information: A preprint of "Unified 3-D structure and projection orientation refinement using Quasi-Newtonian algorithm," by Chao Yang, Esmond G. Ng, and Pawel A. Penczek, can be found at: http://dx.doi.org/10.1016/j.jsb.2004.08.010

— Paul Preuss, LBNL Public Affairs

DataMover Reaches Milestone

(continued from page 1)

data source to stage the files and uses GridFTP to transfer the data. Without DataMover, users would have to manually locate each of the files, then transfer them one by one. Because STAR generates about 400 terabytes of data each year, automating the transfer is critical. The use of GridFTP with large windows, as well as staging, transferring, and archiving multiple files concurrently enable effective end-to-end transfer rates.

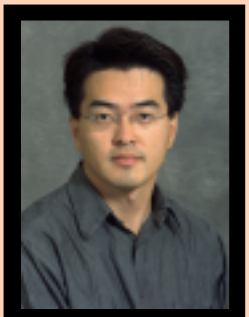

Alexander Sim

"A nice milestone and one more step in the path to success for the LBNL/SRM team," said Brookhaven's Jerome Lauret, who provides server support for the STAR collaboration. "I have to renew my thanks and gratitude for such a tool: it has made our lives so much easier in STAR for the past few years. Thanks and congratulations on a successful demo."

The DataMover also automatically addresses "transient failures," such as failed network connections or problems in a storage system at either end, by automatically retrying until the connection is re-established. File tracking logs also help users monitor problems network slowdowns, files transfers and bottlenecks. In addition, the massive transfer operation also records the files in a file catalog in the target NERSC site.

"In the past, users had to baby-sit all the transfers," said Alex Sim, a member of the Scientific Data Management Research Group and one of the developers of the application. "This application automates all those processes." In addition to Sim, various components of this system were developed by Junmin Gu and Vijaya Natarajan under the leadership of Arie Shoshani, who is the PI for the Storage Resource Management project at LBNL.

The DataMover can work with any storage system that supports an SRM middleware interface. The DataMover is also currently being used by the Earth Systems Grid climate research collaboration between Berkeley Lab, Oak Ridge and Lawrence Livermore national labs, and the National Center for Atmospheric Research. Recently DataMover transferred 4,224 files containing 770 gigabytes from the HPSS at Oak Ridge to NCAR's specialized Mass Storage System. For this purpose the HPSS-HRM was adapted to work with NCAR's MSS.

A short presentation on the SC2004 project can be read at http://sdm.lbl.gov/sc2004/.